Buffer Overflow Analysis for C

Uday P. Khedker

Department of Computer Science and Engineering
Indian Institute of Technology Bombay
Email: uday@cse.iitb.ac.in

arXiv:1412.5400v4 [cs.PL] 28 Dec 2014

Abstract

Buffer overflow detection and mitigation for C programs has been an important concern for a long time. This paper defines a string buffer overflow analysis for C. The key ideas of our formulation are (a) separating buffers from the pointers that point to them, (b) modelling buffers in terms of sizes and sets of positions of null characters, and (c) defining stateless functions to compute the sets of null positions and mappings between buffers and pointers.

This exercise has been carried out to test the feasibility of describing such an analysis in terms of lattice valued functions and relations to facilitate automatic construction of an analyser without the user having to write C/C++/Java code. This is facilitated by avoiding stateful formulations because they combine effects through side effects in states raising a natural requirement of C/C++/Java code to be written to describe them. Given the above motivation, the focus of this paper is not to build good static approximations for buffer overflow analysis but to show how given static approximations could be formalized in terms of stateless formulations so that they become amenable to automatic construction of analysers.

1 Introduction

Low level modelling of strings in C and associated unchecked operations potentially lead to the possibility of buffer overflows. Given the possibility of a potentially fraudulent use of these loop holes in C programs, detection and mitigation of buffer overflows is critical [1, 4, 10, 14, 16].

This paper proposes an analysis to discover buffer overflows. The key ideas of our formulation are (a) separating buffers from the pointers that point to them, (b) modelling buffers in terms of sizes and sets of positions of null characters, and (c) defining stateless functions to compute the sets of null positions and mappings between buffers and pointer. The first idea is not new; the novelty of our work lies in modelling the computations of null position sets and an insistence on stateless formulations. As is customary, we present our formulation in an intraprocedural setting. It should be easy to lift it to an interprocedural setting using standard techniques of interprocedural data flow analysis such as the method of value contexts [7, 8, 12].

Our goal is not to device the best possible static approximations for buffer overflow analysis but to show how given static approximations could be formalized to devise a mathematical formulation which can be transcribed into a declarative specification of the analysis so that it becomes amenable to automatic construction of analysers.

The rest of the paper is organized as follows: Section 2 describes the requirements of buffer overflow analysis by defining the core statements for analysis, the soundness criterion, and our assumptions. Section 3 describes our modelling and defines the analysis in terms of lattices and lattice valued functions and relations. For simplicity of exposition, it assumes that a pointer points to a single buffer at a program point. Section 4 shows the analysis of our running example. Section 5 shows how the model can be extended to allow a pointer to point to multiple buffers at a program point. Section 6 briefly describes the related work

to highlight the trends while Section 7 concludes the paper. Appendix A describes how other statements can be modelled in terms of core statements.

2 Requirements of Buffer Overflow Analysis

Our formulation is guided by the following requirements and assumptions.

2.1 Program Model

We assume the following model of programs.

- A buffer is an array of `char` or pointers to `char` storing C-style strings. It could contain multiple strings and hence multiple null characters (`'\0'`). A pointer may point to any location within a buffer.
- The list of core buffer manipulation statements to be considered for this analysis is as follows.
  - Buffer creation: $x = \mathit{malloc}(i)$ and $x = \mathit{calloc}(i)$ where i is a compile time constant, and $\mathit{free}(x)$.
  - Buffer assignment:
    - $x = y$ and $x = y + i$ where i is a compile time constant.
    - $\mathit{memcpy}(x, y, m)$ which copies a m character long block of memory pointed to by y to the memory pointed to by x.
  - Buffer modification:
    - Direct modification. $x[i] = '\backslash 0'$ and $x[i] = 'c'$ where i is a compile time constant and c is a non-null character.
    - Modification through string functions. $\mathit{strcat}(x,y)$ $\mathit{strcpy}(x,y)$ and their length limited versions ($\mathit{strncat}(x,y,m)$ and $\mathit{strncpy}(x,y,m)$).
  - Buffer reading: Any statement using x, $x[i]$, or $*(x+i)$ as an rvalue or calls to $\mathit{strlen}(x)$.
- A program is viewed as a convention control flow graph with each node representing a single statement.

Appendix A explains how other statements are modelled in terms of these statements.

2.2 Soundness Criterion, Required Approximations, and Assumptions

We assume the following soundness criterion: no buffer overflow should go undetected; false positives about buffer overflow can be tolerated.

In order the ensure soundness, we introduce the following approximation: A buffer has a single set of null positions associated with it. If the sets of a buffer differ along different execution paths reaching a program point, we create an approximate buffer such that regardless of the execution path, every string contained in the original buffer is a substring of some string present in the approximate buffer. This approximation may cause some imprecision in that our analysis may consider longer strings than are actually present in the buffer leading to false positives.

This approximation is implemented in the following manner:

- At a given program point, a buffer could have different null position sets along different control flow paths reaching the program point. Hence, at the join points in the program, the null position sets of a buffer reaching along different control flow paths are intersected.
- We assume that the memory allocated using *malloc* does not contain a null character.

- We assume that the values of integer variables appearing in $malloc(i)$, $x[i]$, or $x = y + i$ are known at compile time or a range analysis has been performed. If range information is available, we choose the low limit of the range of i for $malloc(i)$ and the the high limit of the range of i for $x[i]$ and $x = y + i$. If i is not a compile time constant and its range is not known, we assume its value to be ∞.

3 Formulating Buffer Overflow Analysis

In this section, we model buffers and the relations of pointers holding the addresses of buffers. This is then followed by formulating data flow equations that compute them. As is customary, we present our formulation in an intraprocedural setting. It should be easy to lift it to an interprocedural setting using standard techniques of interprocedural data flow analysis such as the method of value contexts [7, 8, 12].

3.1 Assumptions for Simplifying the Formulation

For simplicity of exposition, we make the following assumptions for the purpose of presenting the formulation. Section 5 extends the formulation by relaxing these assumptions.

- At a given program point, in general, a pointer could point to different buffers along different control flow paths reaching the program point. For simplicity of modelling, we assume that a pointer points to a single buffer.

- Under the assumption that range information is available, we ignore all back edges in the program and view it as a directed acyclic graph. This allows us to handle the situation where pointers to a buffer are advanced in a loop and hence at the loop entry, such a pointer points to two different positions in a buffer. If we do not ignore back edges, such common case usage of pointers will lead to over-approximation of null position sets leading to proliferation of false negatives.

  Loops basically contribute to advancement of pointers or increments of indices. Range information captures these effect. Hence we can reduce false negatives by restricting the merge points to those resulting from forward edges in the control flow graph of the program.

Note that a single pointee assumption does not preclude the possibility that a buffer may be pointed to by multiple pointers. Such a situation is easily handled by our formulation. Indeed, our running example of Figure 1 has pointers x, y, and z all pointing to the same buffer b_1.

3.2 Modelling Buffers and Buffer Pointer Relations

We identify a buffer by its allocation site name and record its size and the set of positions in the buffer that hold the null character $'\backslash 0'$.

Let N be the set of nodes in the control flow graph of the program being analysed, I^+ be the set of positive integers (including 0) and P be the set of pointers to buffers. Using these sets, we define $\mathsf{N}_b^\infty = \{b_n \mid n \in \mathsf{N}\} \cup \{b_\infty\}$ as a set of buffer identities, $\mathsf{S}_b = \mathsf{I}^+ \cup \{\infty\}$ as a set of buffer sizes (or offsets into buffers), and $\mathsf{Z}_b = 2^{\mathsf{S}_b}$ as the set of positions of null characters in a buffer. Buffers are described using the following functions.

- $buf = \mathsf{N}_b^\infty \mapsto \mathsf{S}_b \times \mathsf{Z}_b$ maps a buffer identity to buffer information.

  A buffer is identified by the program point of its creation e.g. b_n denotes the buffer created in statement $n \in \mathsf{N}$. Each buffer has exactly one size and exactly one set of null positions at any given program point. buf also contains a fictitious "undefined" buffer $(b_\infty, \infty, \emptyset)$.

  Buffer size ∞ may also be associated with a valid buffer indicating that at that program point, the buffer size is not a compile time constant. This could be either because the buffer has not been created

along some execution path, or the buffer has been created using a size that is not known at compile time, or the memory has been freed along some execution path.[1]

If the null position set is empty, it indicates that a read will cause a buffer overflow. The presence of ∞ in a null position set indicates that a write has already caused buffer overflow.

- $bpt = \mathsf{P} \mapsto \left(\mathsf{N}_b^\infty \times \mathsf{S}_b\right)$ relates a buffer pointer $x \in \mathsf{P}$ to its pointee buffer $b_n \in \mathsf{N}_b^\infty$ and an offset $i \in \mathsf{S}_b$ into the buffer (because a pointer may point to some position in the middle of a buffer).

We assume that both these functions are total functions. This simplifies our formulations.
We use the following notational conventions:

- α ranges over the set $A = 2^{buf}$ (and thus represent buffer mappings).
- β ranges over the set $B = 2^{bpt}$ (and thus represent pointer to buffer mappings).
- b ranges over buffer identities in N_b^∞.
- w, x, y, z range over the set of pointers P.
- i, j, k, l, m range over the set S_b.
- X, Y, Z range over some set (the types of their elements are evident from the context).

For a buffer map α, the extractor functions $\mathit{size}(\alpha, b)$ and $\mathit{nps}(\alpha, b)$ compute the size and the null position set of a given buffer b in α.

$$(b, k, X) \in \alpha \Rightarrow \mathit{size}(\alpha, b) = k \wedge \mathit{nps}(\alpha, b) = X \tag{1}$$

In other words, $\alpha(b) = (\mathit{size}(\alpha, b), \mathit{nps}(\alpha, b))$.

For a buffer pointer map β, the extractor functions $\mathit{pt}(\beta, x)$ and $\mathit{start}(\beta, x, b)$ return the pointee buffer and the start position of x in a given buffer b.

$$(x, b, i) \in \beta \Rightarrow \mathit{pt}(\beta, x) = b \wedge \mathit{start}(\beta, x, b) = i \tag{2}$$

We compute the positions of null characters and start positions using compile time saturated addition of integers that limits the result to the given (compile time) constant k as defined below.

$$\forall i, j \in \mathsf{Z}_b\colon\ i \oplus_k j = \begin{cases} i + j & i + j \leq k \wedge k \neq \infty \\ \infty & \text{otherwise} \end{cases} \tag{3}$$

Running Example

Figure 1 illustrates our modelling. The details of these computations are explained in Section 4.

3.3 Data Flow Equations

This section describes the lattices of data flow values, the data flow equations, and the flow functions used in the data flow equations.

[1] Since a buffer is identified by its program point of creation, known but dissimilar buffer sizes along different execution paths reaching a program point are not possible.

```
/* Initial situation:
   w points to buffer b0
   x points to buffer b1*/

    z = x+6;
/* After this, z points
    to offset 6 in b1   */

    y = x+4;
/* After this, y points
    to offset 4 in b1   */

/* Call 1. No overflow  */
    strcat(z,y);

/* Call 2. b1 overflows */
    strcat(z,y);

/* Call 3. No overflow  */
    strcpy(y,w);
```

(a) Program

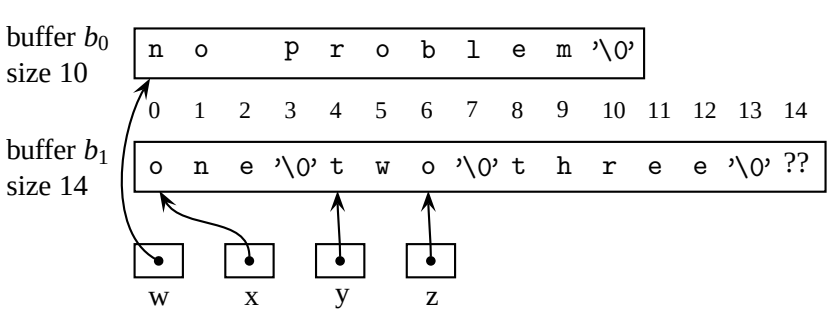

(b) Memory before call 1. "??" indicates garbage value.

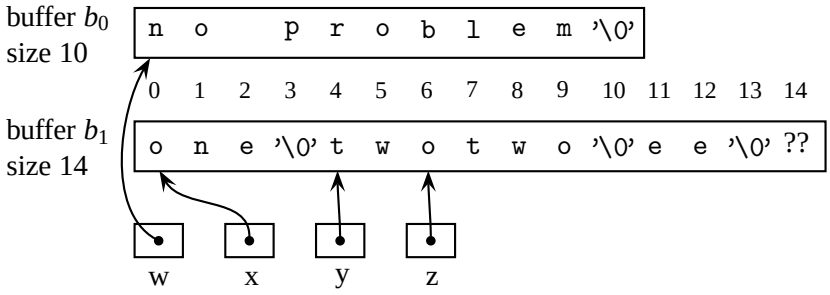

(c) Memory before call 2. "??" indicates garbage value.

Relevant program points	Buffer and pointer mappings	Relevant extractor functions
Before call 1	$\alpha = \{(b_0, 10, \{10\}), (b_1, 14, \{3,7,13\})\}$ $\beta = \{(w, b_0, 0), (x, b_1, 0), (y, b_1, 4), (z, b_1, 6)\}$	$size(\alpha, b_0) = 10, nps(\alpha, b_0) = \{10\}$ $size(\alpha, b_1) = 14, nps(\alpha, b_1) = \{3,7,13\}$ $pt(\beta, w) = b_0,\ start(\beta, w, b_0) = 0$ $pt(\beta, x) = b_1,\ start(\beta, x, b_1) = 0$ $pt(\beta, y) = b_1,\ start(\beta, y, b_1) = 4$ $pt(\beta, z) = b_1,\ start(\beta, z, b_1) = 6$
After call 1, null positions in b_1 change. No other change.	$\alpha = \{(b_0, 10, \{10\}), (b_1, 14, \{3,10,13\})\}$	$size(\alpha, b_1) = 14$ $nps(\alpha, b_1) = \{3,10,13\}$
After call 2, null positions in b_1 change again. No other change.	$\alpha = \{(b_0, 10, \{10\}), (b_1, 14, \{3,\infty\})\}$	$size(\alpha, b_1) = 14$ $nps(\alpha, b_1) = \{3,\infty\}$
After call 3, null positions in b_1 change again. No other change.	$\alpha = \{(b_0, 10, \{10\}), (b_1, 14, \{3,14,\infty\})\}$	$size(\alpha, b_1) = 14$ $nps(\alpha, b_1) = \{3,14,\infty\}$

(d) Modelling buffers and buffer pointers. Presence of ∞ indicates that an overflow has occurred some time.

Figure 1: Example to illustrate modelling of buffers and buffer pointers.

3.3.1 Lattices

Our analysis computes subsets of *buf* and *bpt* flow sensitively at each node $n \in \mathsf{N}$. The lattices of these values are $(A, \sqsubseteq_A)$ and $(B, \sqsubseteq_B\rangle$ where $A = 2^{buf}$ and $B = 2^{bpt}$. The overall lattice L is the product lattice $(A \times B, \sqsubseteq_{AB})$. Its meet operation $\sqcap_{AB}$ is defined in terms of the meet operations $\sqcap_A$ and $\sqcap_B$ of the constituent

lattices A and B.

$$\forall \alpha, \alpha' \in A,\ \forall \beta, \beta' \in B\colon\ (\alpha, \beta) \sqcap_{AB} (\alpha', \beta') = (\alpha \sqcap_A \alpha', \beta \sqcap_B \beta') \tag{4}$$

We first define the meet $\sqcap_A$ for merging the buffers (equations 5 to 7) and then define the meet $\sqcap_B$ for merging the pointer buffer mappings (equation 8).

Merging Buffers

Recall that lattice $(A = 2^{buf}, \sqsubseteq_A)$ where $buf = \mathsf{N}_b^\infty \mapsto \mathsf{S}_b \times \mathsf{Z}_b$ is defined in terms of lattices $(\mathsf{S}_b, \sqsubseteq_S)$ and $(\mathsf{Z}_b, \sqsubseteq_Z)$. Hence $\sqcap_A$ is defined in terms of $\sqcap_S$ and $\sqcap_Z$.

$$\forall \alpha, \alpha' \in A\colon\ \alpha \sqcap_A \alpha' = \left\{ (b,\ ((i \sqcap_S j,\ X \sqcap_Z Y))) \;\middle|\; (b, i, X) \in \alpha, (b, j, Y) \in \alpha' \right\} \tag{5}$$

The definitions of $\sqcap_S$ and $\sqcap_Z$ warrant some explanation.

- The meet $\sqcap_S$ enforces a buffer to have the same size along all paths. In case of inconsistent sizes, the buffer size is recorded as ∞ which indicates that the buffer size is statically undefined (see Section 3.2 for the semantics of ∞).

  $$\forall i, j \in \mathsf{S}_b\colon\ i \sqcap_S j = \begin{cases} i & i = j \\ \infty & \text{otherwise} \end{cases} \tag{6}$$

  The definition of $\sqcap_S$ indicates that ∞ is the $\bot$ element of $(\mathsf{S}_b, \sqsubseteq_S)$. However, it does not contain a $\top$ and hence is only a meet semilattice.

- Given the soundness requirement described in Section 2.2, the $\sqcap_Z$ should approximate the null position sets of buffers by intersecting them. This effectively lengthens the strings in the buffer guaranteeing the soundness of this approximation. However, this idea has a minor flaw: Assume that for a given buffer, we get the set $\{1, 5, 14\}$ along one path and $\{1, 7, \infty\}$ along the other path. We would like to assume that we have a null character at position 1 alone and intersection indeed achieves this. However, we cannot ignore the buffer overflow that has occurred along the second control flow path. An intersection would remove ∞ from the resulting set because it appears in only one set.

  We overcome this problem by observing that a set of null positions X is a subset of $\mathsf{S}_b = \mathsf{I}^+ \cup \{\infty\}$. Hence we define two function $\mathit{fin}(X)$ and $\mathit{inf}(X)$ that partition X into subsets of I^+ and $\{\infty\}$: $\mathit{fin}(X)$ computes the maximal subset of X containing finite numbers, and $\mathit{inf}(X)$ computes the minimal subset of X containing ∞.

  Consider the following sets $X_1 = \{2, 5, 10, \infty\}$, $X_2 = \{2, 5, 10\}$, $X_3 = \{\infty\}$, and $X_4 = \emptyset$. Then,

  - $\mathit{fin}(X_1) = \{2, 5, 10\}$, and $\mathit{inf}(X_1) = \{\infty\}$.
  - $\mathit{fin}(X_2) = \{2, 5, 10\}$, and $\mathit{inf}(X_2) = \emptyset$.
  - $\mathit{fin}(X_3) = \emptyset$, and $\mathit{inf}(X_3) = \{\infty\}$.
  - $\mathit{fin}(X_4) = \mathit{inf}(X_4) = \emptyset$.

  This distinction allows us to use intersection for $\mathit{fin}(X)$ and union for $\mathit{inf}(X)$ in the definition of $\sqcap_Z$.

  $$\forall X, Y \subseteq \mathsf{Z}_b\colon\ X \sqcap_Z Y = (\mathit{fin}(X) \cap \mathit{fin}(Y)) \cup \mathit{inf}(X) \cup \mathit{inf}(Y) \tag{7}$$

Merging Pointer Buffer Mappings

Buffer mappings $\beta, \beta' \in B$ are merged using $\sqcap_B$. Since $B = 2^{bpt}$ where $bpt = \mathsf{P} \mapsto \left(\mathsf{N}_b^\infty \times \mathsf{S}_b\right)$, $\sqcap_B$ is defined in terms of $\sqcap_N$ and $\sqcap_S$:

$$\forall \beta, \beta' \in B\colon\ \beta \sqcap_B \beta' = \left\{ \left(x, b \sqcap_N b', k \sqcap_S k'\right) \;\middle|\; (x,b,k) \in \beta, \left(x,b',k'\right) \in \beta' \right\} \tag{8}$$

The meet operation $\sqcap_S$ is defined in equation (6) earlier; $\sqcap_N$ is similarly defined below.

$$\forall b, b' \in \mathsf{N}_b^\infty\colon\ b \sqcap_N b' = \begin{cases} b & b = b' \\ b_\infty & \text{otherwise} \end{cases} \tag{9}$$

$\sqcap_N$ imposes the restriction that a pointer points to the same buffer along all paths reaching a program point; $\sqcap_S$ ensures that the offsets are also identical.

Since N_b^∞ and S_b are meet semilattice, their product $\mathsf{N}_b^\infty \times \mathsf{S}_b$ is also a meet semilattice. It does not have a $\top$ element and its $\bot$ element is (b_∞, ∞).

3.3.2 Data Flow Equations

Recall that our data flow values are pairs (α, β) that record the buffer mappings and pointer to buffer mappings at each program point.

The data flow equations are as follows:

$$bIn_n = \begin{cases} (\{(b_n, \infty, \emptyset) \mid n \in \mathsf{N}\}, \{(x, b_\infty, \infty) \mid x \in \mathsf{P}\}) & n = \mathit{StartNode} \\ \prod_{p \in pred(n)}^{AB} bOut_p & \text{otherwise} \end{cases} \tag{10}$$

$$bOut_n = \mathit{update_maps}(bIn_n, n) \tag{11}$$

At the start of the program, all buffers are undefined (i.e. their sizes are ∞ and their null pointer sets are $\emptyset$) and all pointers point to undefined buffer (b_∞, ∞). We need this boundary information instead of empty mappings $\emptyset$ because we need to maintain the invariant that our mappings are total functions.

Flow function *update_maps* updates the buffer mappings and pointer to buffer mappings.

$$\mathit{update_maps}(\alpha, \beta, n) = (\mathit{update_buf}(\alpha, \beta, n), \mathit{update_bpt}(\alpha, \beta, n)) \tag{12}$$

For statement n, the destination buffer D_n in buffer mapping α is updated using auxiliary extractor functions $K_n(\alpha, \beta)$ and $R_n(\alpha, \beta)$; the three terms are defined for relevant statements Figure 2 and are explained Section 3.3.3.

$$\mathit{update_buf}(\alpha, \beta, n) = \begin{cases} \alpha\left[D_n \mapsto (K_n(\alpha,\beta), R_n(\alpha,\beta))\right] & \text{Statement } n \text{ involves a string function listed in Figure 2} \\ \alpha & \text{Otherwise} \end{cases} \tag{13}$$

The pointer to buffer maps are updated only for the statements which involve a pointer assignment.

$$\mathit{update_bpt}(\alpha, \beta, n) = \begin{cases} \beta[x \mapsto \{(b_n, 0)\}] & \text{Statement } n \text{ is } x = \mathit{malloc}(k) \text{ or } x = \mathit{calloc}(k) \text{ or } x = \mathit{string}(k) \\ \beta[x \mapsto \beta(y)] & \text{Statement } n \text{ is } x = y \\ \beta[x \mapsto \mathit{shift_offset}(\alpha, \beta, y, i)] & \text{Statement } n \text{ is } x = y + i \\ \beta & \text{Otherwise} \end{cases} \tag{14}$$

where $\mathit{shift_offset}(\alpha, \beta, y, i)$ is used to compute the offset of y in the buffer by shifting it by i.

$$\mathit{shift_offset}(\alpha, \beta, y, i) = \{(b, j \oplus_k i) \mid b = pt(\beta, y), j = \mathit{start}(\beta, y, b), k = \mathit{size}(\alpha, b))\} \tag{15}$$

For brevity, we use the following short forms

- Destination buffer: D_n, New size: $K_n(\alpha,\beta)$, New relevant null position set: $R_n(\alpha,\beta)$
- $pt(\beta,z) = pt_z, size(\alpha,pt_z) = size_z, nps(pt_z) = nps_z, start(\beta,z,pt_z) = start_z, end(\alpha,\beta,z,pt_z) = end_z$

Statement n	D_n	$K_n(\alpha,\beta)$	$R_n(\alpha,\beta)$
$x = malloc(m)$	b_n	m	$\emptyset$
$x = calloc(m)$	b_n	m	$\{i \mid i \leq m\}$
$free(x)$	pt_x	0	nps_x
$x[i] =' \backslash 0'$	pt_x	$size_x$	$nps_x \cup \{i \oplus_{size_x} start_x\}$
$x[i] = c$	pt_x	$size_x$	$nps_x - \{i \oplus_{size_x} start_x\}$
$strcpy\ (x,y)$ $strcat\ (x,y)$ $strncpy\ (x,y,m)$ $strncat\ (x,y,m)$	pt_x	$size_x$	let $length_y = \begin{cases} end_y - start_y & strcpy \text{ or } strcat \\ \min(m+1, end_y - start_y) & strncpy \text{ or } strncat \end{cases}$ $copy_pos_x = \begin{cases} start_x & strcpy \text{ or } strncpy \\ end_x & strcat \text{ or } strncat \end{cases}$ $np_src_{xy} = length_y \oplus_{size_x} copy_pos_x$ $nps_before_{xy} = rnps(nps_x, size_x, 0, start_x, <)$ $np_at_{xy} = \{np_src_{xy}\}$ $nps_after_{xy} = rnps(nps_x, size_x, 0, np_src_{xy}, \geq)$ in $nps_before_{xy} \cup np_at_{xy} \cup nps_after_{xy}$
$memcpy(x,y,m)$	pt_x	$size_x$	let $shift_dist_{xy} = start_x - start_y$ $nps_src_{xy} = rnps(nps_y, size_x, shift_dist_{xy}, start_y, \geq) \cap$ $rnps(nps_y, size_x, shift_dist_{xy}, start_y + m, \leq)$ $nps_before_{xy} = rnps(nps_x, size_x, 0, start_x, <)$ $nps_after_{xy} = rnps(nps_x, size_x, 0, start_x + m, \geq)$ $o_flow_{xy} = \{m \oplus_{size_x} start_x\} - \{m + start_x\} \cup$ $\{m \oplus_{size_y} start_y\} - \{m + start_y\}$ in $nps_before_{xy} \cup nps_after_{xy} \cup nps_src_{xy} \cup o_flow_{xy}$

Figure 2: Buffer mapping extractor functions for statements. We assume that m and i are compile time constants or appropriate limits of ranges (see Section 2.2).

3.3.3 Flow Functions for Computing the Set of Null Positions

The heart of this analysis lies in computing the set of positions of null characters. We use the following functions to achieve this:

1. Functions *size*, *nps*, *pt*, and *start* introduced in equations (1) and (2).
2. Function *rnps* which performs arithmetic on null positions to compute the relevant null positions.

$$rnps(X, sat, shift, pivot, ?) = \{i \oplus_{sat} shift \mid i \in X, i\ ?\ pivot\} \tag{16}$$

It is defined in terms of

- a set of null positions (X),

- a saturation limit (sat),
- an offset ($shift$) in the destination buffer using which the null positions of the source buffer may have to be shifted,
- null positions in the source buffer described in terms of
  - a pivot ($pivot$) around which the null positions are to be examined, and
  - a relational operator ($? \in \{<, \leq, >, \geq, =, \neq\}$) used for comparison with the pivot.

3. Function end which computes the end of a string in a buffer using $rnps$ by finding out the null positions that lie beyond the $start$ of the string and then computing their greatest lower bound (glb_le) under $\leq$ as the partial order. This allows us to restrict the $rnps$ set in a buffer to the first null position after the start position of a given string.

$$end(\alpha, \beta, x, b) = glb_le(rnps(nps(b), size(\alpha, b), 0, start(\beta, x, b), \geq)) \tag{17}$$

Although our primary interest is in computing the minimum among the null positions, glb_le is more convenient. It inherently models the situation when no null position is found because $glb_le(\emptyset) = \infty$ indicating that the string is infinitely long (i.e. the string overflows the buffer).[2]

Set of Null Positions for String Functions

We use the above functions by dividing the null position sets resulting from a string operation into three categories based on the start position s and the end position e of the copied string in the destination buffer. Given x as the destination pointer and y as the source pointer:

1. nps_before. The null positions that remain unchanged in the destination buffer because they appear before s. This is defined by $rnps(nps_x, size_x, 0, start_x, <)$ in Figure 2.

2. np_at. The null position appearing at e. This is defined by np_src in Figure 2. It represents the null position imported from the source buffer.

3. nps_after. The null positions that remain unchanged in the destination buffer because they appear after e. This is defined by $rnps(nps_x, size_x, 0, np_src_{xy}, \geq)$ in Figure 2.

The difference between the $strcat$ and $strcpy$ functions is evident from the definition in Figure 2—for $strcpy$, the starting point of copying is $start_x$ whereas for $strcat$ it is end_x. Further, their length limited versions choose the minimum between the null distance and the length provided.

Set of Null Positions for Memory Copy

We view $memcpy(x, y, m)$ as a special case of $strcpy(x, y)$. The main difference is that $strcpy(x, y)$ copies from y only upto the first null character whereas $memcpy(x, y, m)$ copies m characters from y. Thus the main change is in computing nps_src instead of $np_at = \{np_src\}$. Former includes multiple position whereas the latter computes a single position. Thus we compute

1. nps_src. This identifies all null character positions that lie between $start_y$ and $start_y + m$ in the source buffer. These should be shifted by the start position in the destination buffer. This is computed by the intersection of $rnps(nps_y, size_x, shift_dist_{xy}, start_y, \geq)$ with $rnps(nps_y, size_x, shift_dist_{xy}, start_y + m, \leq)$ where $shift_dist$ represents the distance by which the null positions should be shifted.

2. o_flow. We must include ∞ to indicate overflow if

[2] Given that glb_le computes the greatest number that is smaller than any number in the given set, if the given set does not contain any number, the greatest number that is smaller than no number is ∞.

- $m + start_x > size_x$ (the write operation crosses the destination boundary), or
- $m + start_y > size_y$ (the read operation crosses the source boundary).

This is easily achieved by computing the following set differences $\{m \oplus_{size_x} start_x\} - \{m + start_x\}$ and $\{m \oplus_{size_y} start_y\} - \{m + start_y\}$.

3.4 Overflow Detection

Computation of data flow values $bIn_n/bOut_n$ detects a buffer overflow at a program point at which a buffer is written into by introducing ∞ in the *nps* set. Note that this is a conservative conclusion and may well be a false positive. Our analysis is sound when it concludes the absence of buffer overflow.

For a program statement that merely reads a buffer (eg. call to *strlen*), the $bIn_n/bOut_n$ values remain unmodified. Detecting a potential buffer overflow for a read using a pointer is trivially achieved by

– Checking the buffer identity. If it is b_∞, there is a potential buffer overflow.

– Checking the buffer size. If it is ∞, there is a potential buffer overflow.

– Checking the offset of the pointer. If it is ∞, there is a potential buffer overflow.

– Otherwise, we compute the *end* of the pointer being read. If it is ∞, there is a potential buffer overflow.

4 Running Example Revisited

This section shows the application of our analysis to the example in Figure 1 and describes how buffer overflows can be detected.

4.1 Computing the Set of Null Positions

It is clear from Figure 2 that the only non-obvious computation in this analysis is the computation of R_n. We illustrate it for our running example. The relevant maps and the values of the default extractor functions have already been provided in Figure 1. Observe that for the first two calls, buffer b_1 serves both as source and destination whereas for the third call, buffer b_0 is the source and b_1 is the destination.

- After the first call to $strcat(z,y)$ in the example in Figure 1, we have

$$\begin{aligned}
np_src &= length \oplus_{14} copy_pos \\
&= (end(\alpha,\beta,y,b_1) - start(\beta,y,b_1)) \oplus_{14} end(\alpha,\beta,z,b_1) \\
&= (glb_le(rnps(\{3,7,13\},14,0,4,\geq)) - start(\beta,y,b_1)) \\
&\quad \oplus_{14} \\
&\quad (glb_le(rnps(\{3,7,13\},14,0,6,\geq)) \\
&= (glb_le(\{7 \oplus_{14} 0, 13 \oplus_{14} 0\}) - 4) \oplus_{14} (glb_le(\{7 \oplus_{14} 0, 13 \oplus_{14} 0\})) \\
&= (glb_le(\{7,13\}) - 4) \oplus_{14} (glb_le(\{7,13\})) \\
&= (7-4) \oplus_{14} 7 = 10
\end{aligned}$$

The computation of relevant null positions (i.e. R_n) is as follows:

$$\begin{aligned}
R_n(\alpha,\beta,b_1) &= rnps(\{3,7,13\},14,0,6,<) \cup \{10\} \cup rnps(\{3,7,13\},14,0,10,\geq) \\
&= \{3 \oplus_{14} 0\} \cup \{10\} \cup \{13 \oplus_{14} 0\} \\
&= \{3\} \cup \{10\} \cup \{13\} = \{3,10,13\}
\end{aligned}$$

- After the second call to *strcat* in the example, we have

$$\begin{aligned}
\mathit{np_src} &= \mathit{length} \oplus_{14} \mathit{copy_pos} \\
&= (\mathit{end}(\alpha,\beta,y,b_1) - \mathit{start}(\beta,y,b_1)) \oplus_{14} \mathit{end}(\alpha,\beta,z,b_1) \\
&= (\mathit{glb_le}(\mathit{rnps}(\{3,10,13\},14,0,4,\geq)) - \mathit{start}(\beta,y,b_1)) \\
&\quad \oplus_{14} \\
&\quad (\mathit{glb_le}(\mathit{rnps}(\{3,10,13\},14,0,6,\geq)) \\
&= (\mathit{glb_le}(\{10 \oplus_{14} 0, 13 \oplus_{14} 0\}) - 4) \oplus_{14} (\mathit{glb_le}(\{10 \oplus_{14} 0, 13 \oplus_{14} 0\})) \\
&= (\mathit{glb_le}(\{10,13\}) - 4) \oplus_{14} (\mathit{glb_le}(\{10,13\})) \\
&= (10-4) \oplus_{14} 10 = \infty
\end{aligned}$$

  R_n computation is as follows.

$$\begin{aligned}
R_n(\alpha,\beta,b_1) &= \mathit{rnps}(\{3,7,13\},14,0,6,<) \cup \{\infty\} \cup \mathit{rnps}(\{3,7,13\},14,0,\infty,\geq) \\
&= \{3 \oplus_{14} 0\} \cup \{\infty\} \cup \emptyset \\
&= \{3\} \cup \{\infty\} = \{3,\infty\}
\end{aligned}$$

- The third call $\mathit{strcpy}(y,w)$ involves b_0 as the source buffer and b_1 as the destination buffer.

$$\begin{aligned}
\mathit{np_src} &= \mathit{length} \oplus_{14} \mathit{copy_pos} \\
&= (\mathit{end}(\alpha,\beta,w,b_0) - \mathit{start}(\beta,w,b_0)) \oplus_{14} \mathit{start}(\alpha,\beta,y,b_1) \\
&= (\mathit{glb_le}(\mathit{rnps}(\{10\},10,0,0,\geq)) - 0) \oplus_{14} 4 \\
&= \mathit{glb_le}(\{10 \oplus_{10} 0\}) \oplus_{14} 4 \\
&= \mathit{glb_le}(\{10\}) \oplus_{14} 4 \\
&= 10 \oplus_{14} 4 = 14
\end{aligned}$$

  R_n computation is as follows.

$$\begin{aligned}
R_n(\alpha,\beta,b_1) &= \mathit{rnps}(\{3,\infty\},14,0,4,<) \cup \{14\} \cup \mathit{rnps}(\{3,\infty\},14,0,14,\geq) \\
&= \{3 \oplus_{14} 0\} \cup \{14\} \cup \{\infty\} \\
&= \{3\} \cup \{14\} \cup \{\infty\} = \{3,14,\infty\}
\end{aligned}$$

  Note that the ∞ has not been generated in the computation of *length*, it has been carried forward from the previous *nps* set. In other words, this call does not cause an overflow, the presence of ∞ in R_n indicates that an overflow has occurred in the buffer earlier.

4.2 Overflow Detection

It is clear from Section 3.4 that checking overflow is trivial except in the case where *end* is to be used. We illustrate such uses for our running example in the following cases.

1. Assume that a statement `v = x+14` is added to our running example. Then $\mathit{start}(\beta,v,b_1) = 14$. For simplicity, assume that a call to $\mathit{strlen}(v)$ occurs before the first call to *strcat*. We compute

$$\begin{aligned}
\mathit{end}(\alpha,\beta,v,b_1) &= \mathit{glb_le}(\mathit{rnps}(\{3,7,13\},14,0,14,\geq)) \\
&= \mathit{glb_le}(\emptyset) = \infty
\end{aligned}$$

   Thus $\mathit{strlen}(v)$ may cause a buffer overflow. In this case, this is not a false negative but a certain buffer overflow. However, our analysis does not have any means of distinguishing between a certain overflow and a false negative.

2. For $strlen(y)$ (before the first call to *strcat* in the same example), we compute

$$\begin{aligned} end(\alpha,\beta,y,b_1) &= glb_le(rnps(\{3,7,13\},14,0,4,\geq)) \\ &= glb_le(\{7 \oplus_{14} 0, 13 \oplus_{14} 0\}) \\ &= glb_le(\{7,13\}) = 7 \end{aligned}$$

   The result indicates that the read will encounter a null character in the buffer at position 7 and hence there is no buffer overflow.

3. For $strlen(y)$ (after the first call to *strcat* in the same example), we compute

$$\begin{aligned} end(\alpha,\beta,y,b_1) &= glb_le(rnps(\{3,10,13\},14,0,4,\geq)) \\ &= glb_le(\{10 \oplus_{14} 0, 13 \oplus_{14} 0\}) \\ &= glb_le(\{10,13\}) = 10 \end{aligned}$$

   The result indicates that the read will encounter a null character in the buffer at position 10 and hence there is no buffer overflow.

4. For $strlen(y)$ (after the second call to *strcat* in the same example), we compute

$$\begin{aligned} end(\alpha,\beta,y,b_1) &= glb_le(rnps(\{3,\infty\},14,0,4,\geq)) \\ &= glb_le(\emptyset) = \infty \end{aligned}$$

   The result indicates that the read may not find any null character in the buffer and hence there is a potential buffer overflow. In this case also, this is a certain buffer overflow but we have no way of concluding so.

5 Extensions for Handling Multiple Pointee Buffers of a Pointer

In this section we show how our formulation can be extended to allow a pointer to point to multiple buffers at a program point. This requires a change in the lattice B, in the flow function *update_buf*, and in the computation of extractor function R_n of Figure 2. These changes allow some relaxation in the assumptions about program model.

5.1 Changes in the Lattices

Given $b \neq b'$ and $i \neq j$, we allow the coexistence of triples (x,b,i) and (x,b',j) in a $\beta \in bpt$. However, multiple triples of the kind (x,b,i) and (x,b,j) are still prohibited.

This is achieved by redefining *bpt* for the lattice $(2^{bpt}, \sqsubseteq_B)$ as $bpt = (\mathsf{P} \times \mathsf{N}_b) \mapsto \mathsf{S}_b$. Observe that now we use $\mathsf{N}_b = \{b_i \mid i \in \mathsf{N}\}$ and not N_b^∞ because now we do not need the fictitious "undefined" buffer b_∞. Also, because of the cross product $\mathsf{P} \times \mathsf{N}_b$, $pt(\beta,x)$ now returns a set of buffers rather than a single buffer.

The meet operation $\sqcap_B$ is simplified as follows:

$$\forall \beta,\beta' \in B\colon\ \beta \sqcap_B \beta' = \left\{ \left(x,b,k \sqcap_S k'\right) \;\middle|\; (x,b,k) \in \beta, \left(x,b,k'\right) \in \beta' \right\}$$

Exclusion of b_∞ also leads to a change in the boundary information in the data flow equation for *bIn* in which the triple (x,b_∞,∞) is replaced by the triple (x,b,∞) for *StartNode*.

5.2 Changes in Extractor Function R_n

When we have multiple source destination buffers in a string operation, we use the following approximation by combining the effects of all these buffers: We use the longest string among all destination buffers, the smallest size among all source buffers, and the farthest position of copying among all source buffers. These three changes are reflected in our formulation by computing a single approximate R_n by changing the terms appearing in Figure 2.

We describe the changes in R_n computation for string operations.

- *Computing the longest string among all destination buffers.* We compute length_y separately for each pointee buffer of y and take the largest value.
- *Computing the smallest size among all source buffers.* We take the smallest value of size_x among all pointee buffers of x for computing $\mathit{np_src}_{xy}$.
- *Computing the farthest position of copying among all source buffers.* We take the largest value of $\mathit{copy_pos}_x$ among all pointee buffers of x.
- *Computing R_n.* This involves the following changes.
  - Computing $\mathit{nps_before}_{xy}$ and $\mathit{nps_after}_{xy}$. We take the smallest value of size_x and the largest value of $\mathit{np_src}_{xy}$.
  - Computing $\mathit{np_at}_{xy}$. We take the largest value of $\mathit{np_src}_{xy}$.

Observe that these changes are declarative in the sense that they basically involve ranging over the buffers and computing the maximum or the minimum. Hence a stateless formulation of these changes is easy to write.

Similar changes can be defined for the *memcpy* operation.

5.3 Changes in Flow Function *update_buf*

With the possibility of a pointer pointing to multiple buffers, D_n now becomes a set of buffers and K_n is computed separately for each buffer. However, R_n is common because it is an approximation of all source buffers. A buffer mapping α should accumulate updates in all source buffers in D_n. In other words, when a buffer $b \in D_n$ is updated, the resulting mapping α' should be passed on to *update_buf* for updating some other buffer $b' \in D_n$. This is achieved by defining *update_buf* recursively and passing D_n as the value of argument X for the top level call to *update_buf*.

$$\mathit{update_buf}(\alpha,\beta,n,X) = \begin{cases} \mathit{update_buf}\left(\alpha[b \mapsto (K_n(\alpha,\beta,b),R_n(\alpha,\beta))],\ \mathit{bpt},\ n, X-\{b\}\right) & b \in X \\ \alpha & \text{Otherwise} \end{cases}$$

In each recursive call, we accumulate the updates and the set X becomes smaller. When all updates are accumulated in α, X becomes $\emptyset$.

5.4 Relaxation in Program Model

With the extensions for handling multiple buffers of a pointer at a program point in place, now we do not need to ignore the back edges and we may be able to handle some cases where range information is not available. However, this computation may be expensive as a loop may be iterated many times depending upon the increment in the values of the pointers or array indices.

6 Related Work

Buffer overflow detection and mitigation has been an important concern for a long time. The problem is compounded by the idiosyncrasies of C string operations [15] and non-trivial semantics [3]. There is an abundance of literature on the topic and many tools have been created to address the concern; we cite only a few references as a starting point [1, 4, 10, 14, 16].

A popular approach of addressing this concern has been to perform run time checks using code instrumentation, or expanded representation of pointers and buffers to store bookkeeping information (effected through compilers and changed library function) [2, 6, 13]. The known static analysis methods are characterized by some combination of the following features:

- User annotations in the program [3, 9] in the form of contracts or assertions.
- Range analysis [11, 15] storing the lower and the upper limits of the buffers.
- Integer linear programming [3, 5, 15] to solve constraints to compute the ranges.
- Symbolic computation [11] to store the ranges in terms of expressions rather than in terms of integers.

Most approaches set up constraints flow insensitively (or use flow insensitive pointer analysis) except [11] which does path sensitive analysis by storing relevant path predicates.[3] What is common to all these approaches is that their formulations are stateful algorithms; none of them have a stateless formalization amenable to reasoning and automatic construction of analysers.

In terms of modelling, most approaches do not separate pointers and buffers as we do; they need to explicitly discover aliasing between buffers to record the effects of one change through a pointer, into the buffer of another pointer. In this sense our modelling is closer to the modelling for "intended referent" [6]. However they store the information for run time checking, we use it for static analysis. Besides, unlike them, we do not treat a string as a buffer and instead store a set of null positions within a buffer. Most approaches treat a string as a buffer and hence store a single length with a buffer.

7 Conclusions

Modelling buffers and buffer overflows in terms of

- mappings between pointers and buffers whose addresses they hold,
- buffer information storing size and sets of positions of null characters, and
- defining functions to compute these models

makes it possible to formulate buffer overflow analysis as a data flow analysis of C programs. A key facilitator of this is an emphasis on stateless formalizations of analyses in terms of lattice valued functions and relations. A stateful formulation combines features through side effects recorded in states thereby raising a natural requirement of C/C++/Java code to be written to complement a partially high level specification so that a generator can generate the analyser.

Banishing states from formulations enable higher levels of abstraction. The resulting conciseness, together with higher levels of abstraction, makes the formulations amenable to human reasoning. Further, such formulations allow a generator to check the specifications, combine their features freely, and decide how and where to introduce states in the generated code.

This paper does not claim to define the best buffer overflow analysis; it may well be possible to devise better static approximations. In particular, storing "out of bound" null positions [13] may eliminate many

[3] It avoids a combinatorial explosion by a sparse demand driven computation that uses transitive data and control dependences.

false positives. It may also be interesting to explore the possibility of using symbolic expressions [11] for modelling the set of null character positions.

So long as we have a well defined static approximation, we believe that a stateless mathematical formulation leading to high level specifications amenable to automatic construction of analysers is possible.

Acknowledgements

Dipali Bhutada, Kalyani Zope, Swati Jaiswal, and Vini Kanvar provided useful feedback on the drafts of this paper.

References

[1] C. Cowan, F. Wagle, Calton Pu, S. Beattie, and J. Walpole. Buffer overflows: attacks and defenses for the vulnerability of the decade. In *DARPA Information Survivability Conference and Exposition, 2000. DISCEX '00. Proceedings*, volume 2, 2000.

[2] Dinakar Dhurjati and Vikram Adve. Backwards-compatible array bounds checking for c with very low overhead. In *Proceedings of the 28th International Conference on Software Engineering*, ICSE '06, pages 162–171, New York, NY, USA, 2006. ACM.

[3] Nurit Dor, Michael Rodeh, and Mooly Sagiv. Cssv: Towards a realistic tool for statically detecting all buffer overflows in c. In *Proceedings of the ACM SIGPLAN 2003 Conference on Programming Language Design and Implementation*, PLDI '03, pages 155–167, New York, NY, USA, 2003. ACM.

[4] Pierre-Alian Fayolle and Vincent Glaume. A buffer overflow study: Attacks and defenses. Unpublished manuscipt available at `http://repo.hackerzvoice.net/depot_ouah/bofstd.pdf`. (Last accessed on 16 Dec 2014), 2002.

[5] Vinod Ganapathy, Somesh Jha, David Chandler, David Melski, and David Vitek. Buffer overrun detection using linear programming and static analysis. In *Proceedings of the 10th ACM Conference on Computer and Communications Security*, CCS '03, pages 345–354, New York, NY, USA, 2003. ACM.

[6] Richard W M Jones, Paul H J Kelly, Most C, and Uncaught Errors. Backwards-compatible bounds checking for arrays and pointers in c programs. In *in Distributed Enterprise Applications. HP Labs Tech Report*, pages 255–283, 1997.

[7] U. P. Khedker and B. Karkare. Efficiency, precision, simplicity, and generality in interprocedural data flow analysis: Resurrecting the classical call strings method. In *Proceedings of the International Conference on Compiler Construction*, pages 213–228. Springer-Verlag, 2008.

[8] U. P. Khedker, A. Sanyal, and B. Karkare. *Data Flow Analysis: Theory and Practice*. CRC Press (Taylor and Francis Group), 2009. (Under publication).

[9] David Larochelle and David Evans. Statically detecting likely buffer overflow vulnerabilities. In *Proceedings of the 10th Conference on USENIX Security Symposium - Volume 10*, SSYM'01, Berkeley, CA, USA, 2001. USENIX Association.

[10] Kyung-Suk Lhee and Steve J. Chapin. Buffer overflow and format string overflow vulnerabilities. *Softw. Pract. Exper.*, 33(5):423–460, April 2003.

[11] Lian Li, Cristina Cifuentes, and Nathan Keynes. Practical and effective symbolic analysis for buffer overflow detection. In *Proceedings of the Eighteenth ACM SIGSOFT International Symposium on Foundations of Software Engineering*, FSE '10, pages 317–326, New York, NY, USA, 2010. ACM.

[12] Rohan Padhye and Uday P. Khedker. Interprocedural data flow analysis in soot using value contexts. In *Proceedings of the 2Nd ACM SIGPLAN International Workshop on State Of the Art in Java Program Analysis*, SOAP '13, pages 31–36, New York, NY, USA, 2013. ACM.

[13] Olatunji Ruwase and Monica S. Lam. A practical dynamic buffer overflow detector. In *In Proceedings of the 11th Annual Network and Distributed System Security Symposium*, pages 159–169, 2004.

[14] Hossain Shahriar and Mohammad Zulkernine. Monitoring buffer overflow attacks: A perennial task. *Int. J. Secur. Softw. Eng.*, 1(3):18–40, July 2010.

[15] David Wagner, Jeffrey S. Foster, Eric A. Brewer, and Alexander Aiken. A first step towards automated detection of buffer overrun vulnerabilities. In *In Network and Distributed System Security Symposium*, pages 3–17, 2000.

[16] Misha Zitser, Richard Lippmann, and Tim Leek. Testing static analysis tools using exploitable buffer overflows from open source code. In *Proceedings of the 12th ACM SIGSOFT Twelfth International Symposium on Foundations of Software Engineering*, SIGSOFT '04/FSE-12, pages 97–106, New York, NY, USA, 2004. ACM.

A Modelling Other Statements

We show how most other statements can be modelled in terms of the core statements of our formulation.

- Assignment of string literals (eg. a statement $x =$ "*Hello world*";) is modelled as a sequence of two statements $x = \mathit{malloc}(k); x[k+1] = '\backslash 0'$; where k represents the length of the string literal.
- Reallocation of memory to resize a buffer using $x = \mathit{realloc}(y,k)$ is modelled as a sequence of two statements $x = \mathit{malloc}(k); \mathit{strcpy}(x,y)$.
- Statements $*(x+i) = '\backslash 0'$ and $*(x+i) = 'c'$ are equivalent to $x[i] = '\backslash 0'$ and $x[i] = 'c'$ respectively.
- A declaration of an uninitialized array (eg. *char* $x[k]$;) is modelled as $\mathit{malloc}(k)$ statement.
- A declaration of an initialized array (eg. *char* $x[k] = \{\ldots\}$;) is modelled as a sequence of two statements $x = \mathit{malloc}(k); x[k+1] = '\backslash 0'$; where k represents the number of elements in the array.
- The following string functions are modelled in terms of *strlen* for the purpose of buffer overflow analysis: $\mathit{strcmp}(x,y)$, $\mathit{strncmp}(x,y)$, $\mathit{strchr}(x,y)$, and $\mathit{strrchr}(x,y)$. The function $\mathit{strstr}(x,y)$ finds the first occurrence of string y in string x and is viewed as a combination of $\mathit{strlen}(x)$ and $\mathit{strlen}(y)$ for this analysis. A call to $\mathit{strtok}(x,y)$ is similarly modelled in terms of $\mathit{strlen}(x)$ and $\mathit{strlen}(y)$.

  Some of these functions either return a buffer pointer containing a possible substring (or a collection of substrings). Their lengths are dynamic and hence are not amenable to static analysis. Hence, we view the functions returning pointers to buffers with dynamic lengths, as functions creating buffers of size ∞.
- The remaining memory handling functions can also be similarly modelled and statically approximated in case their result depends on a dynamic value.

We do not rule out the possibility of better static approximations and expect them to be handled in a similar manner provided they can be defined precisely in terms of stateless functions.